

- 1 1. A redundant array of inexpensive disks
2 comprising:
3 a plurality of disks;
4 a bus coupled to said disks and adapted to transmit
5 data blocks from said disks to a receiving device; and
6 a reconstructor adapted to reconstruct a data block
7 of one of said disks with associated data and parity blocks
8 from other of the disks and to transmit the reconstructed
9 block to the receiving device in response to determining
10 that the one of the disks is slowly responding.
- 1 2. The redundant array of claim 1 wherein a slowly
2 responding disk does not complete transmission of a
3 requested data block within a predetermined time.
- 1 3. The redundant array of claim 2, further
2 comprising:
3 a processor adapted to send a signal to the
4 reconstructor in response to the one of the disks slowly
5 responding; and
6 wherein the reconstructor is adapted to transmit the
7 reconstructed block to the receiving device in response to
8 receiving the signal.
- 1 4. The redundant array of claim 3, wherein the
2 processor is adapted to request the other of the disks to
3 send associated data and parity blocks to the reconstructor
4 in response to the one of the disks responding slowly.

Patent 4,093,650

1 5. The redundant data array of claim 3, wherein
2 the processor couples to the plurality of disks and is
3 adapted to control writes of associated data and parity
4 blocks to said disks.

1 6. A method of transmitting data from a redundant
2 array of inexpensive disks (RAID configuration), comprising:
3 requesting a first disk of the RAID configuration to
4 transmit a first data block stored therein to a receiving
5 device;
6 reconstructing the first data block from associated
7 data stored in other disks of the RAID configuration; and
8 transmitting the reconstructed first data block
9 directly to the receiving device; and
10 wherein the reconstructing and the transmitting are
11 performed in response to the first disk providing data in
12 the first data block, but not completing a transmission of
13 the first data block within a predetermined time.

1 7. The method of claim 6, wherein the associated
2 data comprises at least one data block and a parity block.

1 8. The method of claim 6, wherein the
2 reconstructing and the transmitting are performed in
3 response to determining that the first disk is slowly
4 responding.

1 9. The method of claim 6, further comprising:
2 providing a time signal at a predetermined time after the
3 requesting; and

4 wherein the transmitting is in response to an
5 occurrence of the time signal before the first disk
6 completes transmission of the first data block.

1 10. The method of claim 7, wherein the
2 reconstructing comprises calculating a bit-by-bit exclusive-
3 OR for corresponding bits of data and parity blocks
4 associated with the first data block.

1 11. A redundant array of inexpensive disks (RAID)
2 configuration to transmit data blocks to a receiving device,
3 comprising:

4 a plurality of disks adapted to store the blocks and
5 to transmit the stored blocks to the receiving device;

6 a processor to control reads from and writes to the
7 plurality of disks and to determine if one of the disks is
8 slowly responding; and

9 a device adapted to reconstruct a block stored in a
10 slowly responding one of the disks from associated blocks
11 stored in others of the disks in response to the processor
12 determining that the one of the disks is slowly responding.

1 12. The RAID of claim 11, wherein the processor is
2 configured to determine that the one of the disks is slowly
3 responding if the one disk does not finish transmitting a
4 requested data block within a predetermined time.

0995504-12001

1 13. The RAID configuration of claim 12, further
2 comprising:

3 a timer coupled to the processor and adapted to
4 count the predetermined time in response to the processor
5 sending a request to a disk to transmit a data block.

1 14. A two-level redundant array of inexpensive
2 disks (RAID), comprising:

3 a first level processor; and

4 a plurality of second level redundant arrays of
5 inexpensive disks, the first level processor adapted to
6 write first level blocks to and read first level blocks from
7 the second level arrays; and

8 each second level array comprising:

9 a plurality of disks adapted to store second level
10 blocks and to retrieve stored blocks; and

11 a first level device to reconstruct a particular
12 first level block from associated first level blocks and to
13 transmit the reconstructed first level block to a receiving
14 device in response to a second level array responding
15 slowly.

1 15. The RAID configuration of claim 14, wherein
2 the first level processor is adapted to signal the first
3 level device that one of the second level arrays is
4 responding slowly if the one of the second level arrays does
5 not complete a transmission of a requested first level block
6 in a predetermined time.

1 16. The RAID configuration of claim 14, each
2 second level array further comprising:

3 a second level processor to request reads of and
4 writes to the disks of the associated second level array.